

Stochastic Atmosphere Modeling for Risk Adverse Aerocapture Guidance

Jack Ridderhof and Panagiotis Tsiotras
Georgia Institute of Technology
270 Ferst Dr.
Atlanta, GA 30313
{jridderhof3, tsiotras}@gatech.edu

Abstract—In this paper we consider the problem of aerocapture guidance inside an atmosphere with the density randomly perturbed according to a Gaussian process model. We formulate the problem of solving for a single bank event as a stochastic optimization problem and then, in terms of a 3-DoF Martian aerocapture example, we study the properties of the proposed approach compared to a standard deterministic approach. The stochastic approach is shown to decrease the risk that a large delta- V is required for maneuvering after atmospheric flight to reach a target apoapsis and periapsis.

TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. PROBLEM FORMULATION.....	2
3. DENSITY PERTURBATION MODEL	3
4. AEROCAPTURE MISSION DESIGN.....	4
5. NUMERICAL EXAMPLE.....	5
6. CONCLUSION	6
ACKNOWLEDGMENTS	6
REFERENCES	6
BIOGRAPHY	7

1. INTRODUCTION

Aerocapture is a promising technology for increasing the usable payload mass for planetary science missions, but the benefits come at the cost of increased complexity and uncertainty in terms of performance. Despite these drawbacks, aerocapture has been extensively studied for use in planetary missions and is considered to be an enabling technology for an orbiter to explore the ice giants [1], [2]. Other proposed uses of aerocapture include increasing the mass delivered to Mars in support of human missions and enabling a smallsat Venus orbiter [3], [4], [5]. While, as of the writing of this paper, aerocapture has not yet been demonstrated in an actual planetary mission, orbit adjustments using aerobreaking have been successfully demonstrated at Earth, Venus, and Mars [6].

An aerocapture mission involves atmospheric flight around a planet for which the atmosphere may not be completely characterized. If the true atmospheric properties do not exactly match the preflight prediction, then the orbit of the spacecraft following atmospheric flight cannot be predetermined with high accuracy and confidence. Indeed, the distribution of possible orbits following atmospheric flight is a function of both the variations in the true atmosphere from the preflight predictions and the state of the spacecraft during the flight. By using feedback control, an aerocapture guidance system

is instrumental in shaping the distribution of atmosphere exit conditions so that the probability of failure is acceptably low.

In this work we model the density profile of an atmosphere as a random process with known statistics, and study the effect of closed-loop guidance on the resulting probability density of atmospheric exit conditions. Studying aerocapture guidance in this framework offers several insights. By considering the spacecraft state to be a random variable, we may define an optimal aerocapture maneuver in terms of the shape of the resulting probability distribution [7], [8]. This permits a system designer to minimize the risk associated with a random outcome compared to simply minimizing a mean value. In particular, we penalize the Conditional Value-at-Risk (CVaR) of delta- V required for maneuvering after atmospheric flight since CVaR is a coherent measure of risk that is well suited to optimization [9], [10].

It is well known that optimal control for linear stochastic systems with quadratic cost and additive Gaussian noise may be obtained by replacing the stochastic terms with their expected values and solving the resulting deterministic system [11]. This assumption is referred to as certainty equivalence since the controls are obtained by assuming random variables to be equivalent to their expected values [12]. However, in general, the optimal control for a nonlinear stochastic system depends on the distribution of the state rather than just the mean [13]. For aerocapture, given the spacecraft state at some time, the density profile of the atmosphere must be known in order to accurately predict the spacecraft state at a future time. Therefore, the state of the aerocapture system includes both the spacecraft and atmosphere states. It follows that, since aerocapture dynamics are nonlinear, the optimal aerocapture guidance commands should be computed as a function of the distribution of the stochastic atmosphere state.

Computing guidance commands as a function of the distribution of the atmospheric density process presents new challenges for navigation. In order to achieve closed-loop control, the spacecraft must successively, as new measurements become available, compute the updated conditional distribution of the atmospheric density profile. In this paper, we use standard results from Gaussian process theory to compute the conditional distribution, with respect to past measurements, of the atmospheric density process in terms of its conditional mean and covariance. The resulting navigation solution resembles a Kalman filter with the atmosphere as a state [14]. We then, using Monte Carlo, numerically integrate the dynamics with respect to sampled atmospheric density profiles to approximate the distribution of future spacecraft states induced by the conditional atmosphere distribution.




Figure 1. α -VaR and α -CVaR for a normal random variable, where the distribution function is shown on the left and the density function is shown on the right.

Notation and Preliminaries

For functions $f : \mathbb{R} \rightarrow \mathbb{R}$, $g : \mathbb{R}^2 \rightarrow \mathbb{R}$ and vectors $x = (x_1, \dots, x_n)$, $y = (y_1, \dots, y_m)$, we write $f(x)$ to denote the vector with coordinates $f(x_i)$ and $g(x, y)$ to denote the matrix with entries $g(x_i, y_j)$.

For a random variable z , the α -Value-at-Risk (α -VaR) of z is the smallest value β such that the probability of $z \leq \beta$ is greater than α , that is,

$$\alpha\text{-VaR}(z) = \min\{\beta \in \mathbb{R} : \mathbb{P}(z \leq \beta) \geq \alpha\} \quad (1)$$

for $\alpha \in [0, 1)$. In other words, the α -VaR is the minimum threshold for the upper α percentile of z . While the term value at risk originated in finance [15], a similar concept is used for what are called chance constraints in stochastic optimal control (e.g., the probability that the robot hits a wall should be less than a prescribed value) [16], [7].

The α -Conditional VaR (α -CVaR) [15] of the random variable z is the expected value of z conditioned on the event that z is greater than its α -VaR:

$$\alpha\text{-CVaR}(z) = \mathbb{E}[z | z \geq \alpha\text{-VaR}(z)], \quad (2)$$

that is, $\alpha\text{-CVaR}(z)$ is the mean of the upper α percentile of z (see Figure 1).

2. PROBLEM FORMULATION

Consider a spacecraft in atmospheric flight around a spherical body trimmed for constant lift. The longitudinal equations of motion for this vehicle are given by [17]

$$\dot{r} = v \sin \gamma, \quad (3a)$$

$$\dot{v} = -\frac{\rho v^2}{2\beta} - \frac{\mu \sin \gamma}{r^2}, \quad (3b)$$

$$\dot{\gamma} = \frac{\rho v (C_L/C_D)}{2\beta} \cos \sigma - \left(\frac{\mu}{r^2} - \frac{v^2}{r} \right) \frac{\cos \gamma}{v}, \quad (3c)$$

where r is the radial position of the spacecraft, v is the spacecraft planet-relative velocity, γ is the spacecraft planet-relative flight path angle defined positive up, μ is the gravitational parameter of the planet, ρ is atmospheric density, and $\beta = m/SC_D$ is the spacecraft ballistic coefficient in terms of mass m , reference area S , and drag coefficient C_D . The in-plane component of the lift is controlled by rolling the vehicle




Figure 2. Atmospheric flight segment with lift L , drag D , gravitational acceleration g , and control $u = \cos \sigma$.

about the velocity vector by a bank angle σ . We assume that the response time to roll commands is sufficiently small so that we may consider the control input to be the vertical fraction of the lift magnitude, given by

$$u = \cos \sigma. \quad (4)$$

Define the atmosphere interface (AI) to be the time t_i when $r(t_i) = r_{\text{atm}}$, where $r_{\text{atm}} > 0$ is a fixed and known radial distance marking the edge of the meaningful atmosphere. At AI we assume that the spacecraft states are known to be

$$v(t_i) = v_i, \quad \gamma(t_i) = \gamma_i < 0. \quad (5)$$

Let the exit time t_e be defined as the first time after t_i where the spacecraft radial position is equal to r_{atm} . Flight between times t_i and t_e will be referred to as atmospheric flight. See Figures 2 and 4.

The density of the atmosphere $\rho(h)$ is assumed to be a random process as a function of altitude $h = r - R_p$ defined over the altitude range $[0, h_{\text{atm}}]$, where $h_{\text{atm}} = r_{\text{atm}} - R_p$ with R_p the planet radius. It follows that the spacecraft state (r, v, γ) is a random vector for all times $t \geq t_i$. We assume that the spacecraft has perfect knowledge of its own state during flight. On the other hand, the density of the atmosphere is not known and must be estimated from noisy measurements y_k taken at times t_k according to the model

$$y_k = \delta \rho[h(t_k)] + \eta_k, \quad (6)$$

where η_k are independent and identically distributed zero-mean Gaussian random variables with variance σ_{η}^2 , and where

$$\delta \rho = \rho / \bar{\rho} - 1 \quad (7)$$

is the density variation from a known mean density $\bar{\rho} : [0, h_{\text{atm}}] \rightarrow (0, \infty)$. In terms of the vectors Y^k and H^k of the first k measurements and measurement altitudes given by

$$Y^k = (y_1, \dots, y_k), \quad H^k = (h_1, \dots, h_k), \quad (8)$$

where $h_i = h(t_i)$, let $\mathcal{I}_k = (Y^k, H^k)$ denote the information of the first k measurements. Note that we are assuming that the altitudes $h(t_i)$ at which measurements are taken are known exactly.

Suppose for a moment that the density is a known exponential function. Then the control that minimizes the delta- V required to reach a target periapsis and apoapsis is composed

of a single bank maneuver from full lift up to full lift down [18], that is,

$$u(t; t_s) = \begin{cases} +1, & t < t_s, \\ -1, & t \geq t_s, \end{cases} \quad (9)$$

where the time t_s is called the switching time. It follows that solving for the optimal control reduces to a line search problem, and therefore the optimal control may be continuously resolved onboard during flight. However, if there are perturbations in the atmospheric density, then the guidance must make corrections which are not accounted for in this control structure. A common solution is to solve for the switch time under the assumption that $u(t > t_s) > -1$, which has the effect of reserving future control authority to make corrections [18]. Regardless, in this work we assume that the control is of the form (9) so the optimal control problem reduces to a single dimensional optimization problem.

Finally, before we pose the problem of solving for an optimal switch time, we must define the optimization objective. If there are no atmospheric density perturbations, then clearly the optimal control during atmospheric flight should be chosen to minimize fuel, or, equivalently, the delta- V , required to reach the final desired orbit. This objective has been shown to be nearly equivalent to minimizing the apoapsis targeting error, which is an analytically more tractable objective [19]. Therefore, modern algorithms solve for the switch time by numerically predicting the apoapsis as a function of the current estimated density profile [18], [5]. These algorithms, and numerical-predictor corrector (NPC) algorithms in general, implicitly rely on the heuristic that successively, as new measurements become available, minimizing the predicted delta- V will, in turn, reduce the risk of mission failure caused by not having enough fuel to complete the required post-atmospheric flight maneuvers. That is, these algorithms assume certainty equivalence.

In this work, we explicitly model the atmospheric density perturbations, and so we may directly minimize a measure of risk. Specifically, at a time t_g after which k measurements have been collected, and for a given $\alpha \in [0, 1]$, we define the optimal switch time to be the time $t_s^* \geq t_g$ that minimizes the cost

$$J(t_s; \mathcal{J}_k) = \alpha - \text{CVaR} [\Delta V(t_s) | \mathcal{J}_k], \quad (10)$$

where $\Delta V(t_s)$ is a function of the dynamics (3), the density variation model, and a fixed post-atmospheric flight maneuver plan uniquely determined by the spacecraft state at atmosphere exit. In other words, the optimal switch time minimizes the expected upper α percentile of delta- V based on the information of all proceeding measurements. Note that setting α to zero reduces the problem to minimizing the expected delta- V .

3. DENSITY PERTURBATION MODEL

We model the atmospheric density variation as a zero-mean Gaussian random process $h \mapsto \delta\rho(h)$ with a known covariance function $\Sigma : [0, h_{\text{atm}}]^2 \rightarrow \mathbb{R}$. In particular, we model the atmospheric density variations as an Ornstein–Uhlenbeck process with altitude-dependent variance based on the MarsGRAM perturbation model [20].




Figure 3. Sample density variation functions from MarsGRAM (left) and from the Gaussian process model (right). 2σ deviations are shown by shaded regions.

That is, for $h_1, h_2 \in [0, h_{\text{atm}}]$, we assume

$$\begin{aligned} \Sigma(h_1, h_2) &= \exp\left(-\frac{|h_1 - h_2|}{h_{\text{scale}}}\right) \\ &\times \begin{cases} b(\min(h_1, h_2)), & \min(h_1, h_2) < h_{\text{trans}}, \\ \sigma_{\rho, \text{max}}^2, & \min(h_1, h_2) \geq h_{\text{trans}}, \end{cases} \end{aligned} \quad (11)$$

where h_{scale} is the scale height, and where

$$b(h) = \sigma_{\rho, \text{max}}^2 \exp\left(\frac{h - h_{\text{trans}}}{c_{\text{scale}}}\right). \quad (12)$$

The constants h_{trans} and c_{scale} determine the scale of the exponential variance model and $\sigma_{\rho, \text{max}}^2$ is the maximum density variance, which is realized for altitudes $h \geq h_{\text{trans}}$. This model is compared with MarsGRAM in Figure 3, which shows very close correlation between the two models.

For onboard guidance and navigation calculations, we will only consider a finite collection of altitude values

$$\vec{h} = (h_{p_1}, \dots, h_{p_n}) \quad (13)$$

and consider the corresponding atmospheric density variation at these altitudes:

$$\delta\vec{\rho} = \delta\rho(\vec{h}) = (\delta\rho(h_{p_1}), \dots, \delta\rho(h_{p_n})). \quad (14)$$

The vector $\delta\vec{\rho}$ is thus a finite-dimensional approximation of the atmosphere state that can be represented in the navigation filter. It follows that the atmospheric density at the altitudes $\vec{h}$ is

$$\vec{\rho} = \bar{\rho}(\vec{h}) [1 + \delta\vec{\rho}]. \quad (15)$$

At the start of atmospheric flight, before any measurements are taken, $\delta\vec{\rho}$ is a zero-mean Gaussian random vector with covariance matrix $\Sigma(\vec{h}, \vec{h})$. This is analogous to the state estimate in a Kalman filter prior to taking measurements. Since we have assumed that atmospheric density variation measurements have additive Gaussian noise according to

(6), the density variation vector conditioned on measurement information is a Gaussian vector

$$\delta\vec{\rho}|\mathcal{J}_k \sim \mathcal{N}(\mu(k), Q(k)) \quad (16)$$

with mean

$$\begin{aligned} \mu(k) &= \mathbb{E}(\delta\vec{\rho}|\mathcal{J}_k) \\ &= \Sigma(\vec{h}, H^k) [\Sigma(H^k, H^k) + \sigma_\eta^2 I]^{-1} Y^k \end{aligned} \quad (17)$$

and covariance

$$\begin{aligned} Q(k) &= \text{cov}(\delta\vec{\rho}|\mathcal{J}_k) = \Sigma(\vec{h}, \vec{h}) \\ &\quad - \Sigma(\vec{h}, H^k) [\Sigma(H^k, H^k) + \sigma_\eta^2 I]^{-1} \Sigma(\vec{h}, H^k). \end{aligned} \quad (18)$$

We draw a fixed number N_s of i.i.d. samples $(\delta\vec{\rho}_i)_{i=1}^{N_s}$ from the distribution $\mathcal{N}(\mu(k), Q(k))$ and compute the corresponding samples

$$\Delta V_i = \Delta V_i(\delta\vec{\rho}_i, t_s), \quad i = 1, \dots, N_s \quad (19)$$

by integrating the dynamics for a fixed switch time t_s . The α -VaR and α -CVaR of ΔV are then approximated using the samples (ΔV_i) as follows. First, we order the sample values so that $\Delta V_i \leq \Delta V_j$ when $i \leq j$. The α -CVaR is then approximated by

$$\alpha\text{-CVaR}[\Delta V(t_s)|\mathcal{J}_k] \approx \frac{1}{N_s - N_{\text{VaR}} + 1} \sum_{i=N_{\text{VaR}}}^{N_s} \Delta V_i, \quad (20)$$

where $N_{\text{VaR}} = \min\{i \in \{1, \dots, N_s\} : i/N_s \geq \alpha\}$.

4. AEROCAPTURE MISSION DESIGN

We consider an aerocapture maneuver designed to minimize the fuel required, via a sequence of impulsive maneuvers after atmospheric flight, to reach a target apoapsis and periapsis. See Refs. [19], [4], [21]. The mission is divided into five segments, as shown in Figure 4, with the following transitions:

- I-II Start atmospheric flight
- II-III End atmospheric flight; lower apoapsis if necessary
- III-IV Raise periapsis to target
- IV-V Correct apoapsis to target

The spacecraft's orbit immediately following the atmospheric flight segment is determined by the spacecraft state at the time of atmospheric exit t_e , denoted by (r_e, v_e, γ_e) . In particular, we fix functions describing the impulsive maneuvers in terms of the atmospheric exit conditions so that the total ΔV required to reach the target orbit, described by the target periapsis r_{pt} and target apoapsis r_{at} , which are shown in Figure 4, is uniquely determined by the spacecraft state at time t_e . In the following, we derive the expression for ΔV as a function of the vehicle states at atmospheric exit.

Exit Correction Burn 1

After leaving the atmosphere, the spacecraft is not necessarily in a closed orbit about the planet, and therefore the spacecraft may need to perform a corrective maneuver to be captured into a closed orbit. Indeed, if the apoapsis following the atmospheric exit is greater than a prescribed maximum value $r_{a1, \max}$, then a correction burn is made so that the apoapsis is equal to $r_{a1, \max}$. To minimize the fuel cost, this burn should




Figure 4. Aerocapture mission overview.

be performed as close as possible to the periapsis, which is directly after atmospheric exit.

The orbit following atmospheric exit has apoapsis

$$r_{a1} = a_{\text{III}} \left(1 + \sqrt{1 - \frac{r_{\text{atm}}^2 v_e^2 \cos^2 \gamma_e}{\mu a_{\text{III}}}} \right), \quad (21)$$

where a_{III} is the semi-major axis given by

$$a_{\text{III}} = \frac{\mu}{2\mu/r_{\text{atm}} - v_e^2}. \quad (22)$$

Denoting the velocity before and after a maneuver with superscript $+$ and $-$, respectively, the first correction burn is therefore given by

$$\Delta V_1 = \begin{cases} v_e^- - v_e^+, & r_{a1} > r_{a1, \max}, \\ 0, & \text{otherwise,} \end{cases} \quad (23)$$

where the thrust is directed opposite the velocity vector, and where, if $r_{a1} > r_{a1, \max}$,

$$v_e^+ = \sqrt{\frac{2\mu r_{a1, \max}(r_{a1, \max} - r_{\text{atm}})}{r_{\text{atm}}(r_{a1, \max}^2 - r_{\text{atm}}^2 \cos^2 \gamma_e)}}, \quad (24)$$

otherwise $v_e^+ = v_e^-$. In the following analysis we denote $v_e = v_e^+$.

Periapsis Raising Burn 2

At the first apoapsis, the periapsis is raised from the current periapsis r_{p1} to the prescribed target periapsis r_{pt} . The velocity before this correction burn is given by

$$v_{a1}^- = \left(v_e^2 - 2\mu \left(\frac{1}{r_{\text{atm}}} - \frac{1}{r_{a1}} \right) \right)^{1/2}, \quad (25)$$

where r_{a_1} as in (21). Similarly, the velocity after this correction burn is given by

$$v_{a_1}^+ = \sqrt{2\mu} \left(\frac{1}{r_{a_1}} - \frac{1}{2(r_{a_1} + r_{p_t})} \right)^{1/2}, \quad (26)$$

and therefore

$$\Delta V_2 = |v_{a_1}^+ - v_{a_1}^-|. \quad (27)$$

Apoapsis Correction Burn 3

At the first periapsis following the atmospheric flight segment, the apoapsis is corrected from the current value r_{a_1} to the prescribed target value r_{a_t} . The velocity before this maneuver is

$$v_{p_t}^- = \left((v_{a_1}^+)^2 - 2\mu \left(\frac{1}{r_{a_1}} - \frac{1}{r_{p_t}} \right) \right)^{1/2}, \quad (28)$$

and the target velocity after the maneuver is

$$v_{p_t}^+ = \sqrt{2\mu} \left(\frac{1}{r_{p_t}} - \frac{1}{2(r_{p_t} + r_{a_t})} \right)^{1/2}, \quad (29)$$

and therefore

$$\Delta V_3 = |v_{p_t}^+ - v_{p_t}^-|. \quad (30)$$

Finally, the total cost is the sum

$$\Delta V = \Delta V_1 + \Delta V_2 + \Delta V_3. \quad (31)$$

5. NUMERICAL EXAMPLE

Consider an aerocapture at Mars with a fixed true density variation $\delta\rho$ provided by MarsGRAM and with parameters given in Table 1. For a moment, suppose that the switch time is fixed to be 165 seconds after AI, which results in the trajectory shown in Figure 5. The spacecraft does not know the true density profile and only has access to noisy measurements according to the model (6) with measurements taken at 1 Hz. The density profile $\delta\bar{\rho}$ is evaluated at altitudes $\bar{h} = (31, 32, \dots, 224)$ km.

The conditional distribution of the atmospheric density variation evaluated at 165 seconds after AI is shown in Figure 6. This example demonstrates that if the measurement variance is sufficiently large, then the density is not exactly known even for altitudes at which measurements have been taken. The corresponding conditional distribution of delta- V is approximated by $N_s = 2000$ samples from the conditional density variation. As the spacecraft altitude decreases and more measurements are collected, the conditional distribution of delta- V approaches its true value. This is shown in Figure 7 for times $t \in t_i + \{0, 165, 180, 220, 240\}$. At each guidance update time, we can see that fixing a switch time determines a distribution delta- V . Optimizing over the switch time in this setting equates to shaping the conditional delta- V distribution in some desirable way. In contrast to an approach where only a single estimate of the atmosphere is used for the optimization, here we are able to optimize with respect to the best and worse case atmosphere models.

Now suppose that at 160 seconds after AI the guidance system performs an update of the switch time. The measurements collected up until that time inform a conditional

Table 1. Simulation parameters

Parameter	Value	Units
R_p	3.397e3	km
h_{atm}	220	km
μ	4.282e13	m ³ /s ²
$r(t_0)$	$300 + R_p$	km
$v(t_0)$	6.1e3	km/s
$\gamma(t_0)$	-17.7	deg
β	150	kg/m ²
C_L/C_D	0.18	-
N_s	2000	-
h_{scale}	11.1	km
$\sigma_{\rho, \text{max}}^2$	1480	(kg/m ³) ²
h_{trans}	120	km
c_{scale}	20	km




Figure 5. Sample trajectory with fixed $t_s = 165$. Cross marks correspond to epochs shown as rows in Figure 7. The switch time is shown by the dotted line.

distribution of the density variation profile, which, in turn, determines a conditional distribution of delta- V . Figure 8 shows, for different switch times, the α -CVaR of this distribution for $\alpha \in \{0.95, 0.99\}$ together with the certainty equivalence and true costs. The certainty equivalence cost is derived by assuming that the conditional mean density is the true density, and the true cost is solved by knowing the true density profile, which is not known to the spacecraft. While this comparison is for a single true atmosphere sample, Figure 8 nonetheless demonstrates that the CVaR cost is not equal to the certainty equivalence cost and that the minimizing switch time depends on α . The reason for this difference can be seen in the conditional delta- V probability density plots for two different switch times, which are given in Figure 9. In this example, we see that the later switch time slightly decreases the probability of $\Delta V \geq 600$ m/s despite having a higher mean.




Figure 6. Sample state of the onboard density estimation at 165 seconds after AI.

6. CONCLUSION

Aerocapture trajectories are highly sensitive to variations in the atmospheric density of the planet, which must be accounted for when designing an aerocapture guidance system. This problem has been addressed in this paper by explicitly modeling the variations in atmospheric density as a Gaussian random process so that the sensitivity to density perturbations may be studied statistically. It was shown that, in this framework, aerocapture guidance commands should be computed as a function of the distribution of atmospheric density profiles conditioned on in-situ measurements. The aerocapture guidance problem was then posed as a stochastic optimization problem with the objective of minimizing risk of mission failure, which is measured in terms of the CVaR of ΔV required to reach a target apoapsis and periapsis. The resulting guidance scheme was shown, via a numerical example, to decrease the risk of a large required ΔV when compared to a deterministic method.

ACKNOWLEDGMENTS

This work was supported by a NASA Space Technology Research Fellowship.

REFERENCES

- [1] M. K. Lockwood, K. T. Edquist, B. R. Starr, B. R. Hollis, G. A. Hrinda, R. W. Bailey, J. L. Hall, T. R. Spilker, M. A. Noca, and N. O’Kongo, “Aerocapture systems analysis for a Neptune mission,” NASA, Tech. Rep. TM-2006-214300, 2006.
- [2] R. G. Deshmukh, S. Dutta, and D. A. Spencer, “Investigation of direct force control for planetary aerocapture at Neptune,” in *29th AAS/AIAA Space Flight Mechanics Meeting*, no. AAS 19-212, Maui, Hawaii, January 2019.
- [3] R. D. Braun and R. W. Powell, “Predictor-corrector




Figure 7. Sample evolution of the density estimation and the corresponding distribution of the conditional ΔV . From the top, the times shown are 0, 165, 180, 220, and 240 seconds after AI. The true atmosphere is shown as a solid black line and the current altitude is shown by a dashed line.

guidance algorithm for use in high-energy aerobraking system studies,” *Journal of Guidance, Control, and Dynamics*, vol. 15, no. 3, pp. 672–678, 1992.

- [4] R. W. Powell and R. D. Braun, “Six-degree-of-freedom guidance and control analysis of Mars aerocapture,” *Journal of Guidance, Control, and Dynamics*, vol. 16, no. 6, pp. 1038–1044, November–December 1993.
- [5] E. Roelke, M. Werner, and R. D. Braun, “Single-stage drag modulation GNC performance for Venus aerocapture demonstration,” in *AIAA Scitech 2019 Forum*, 2019.
- [6] D. A. Spencer and R. Tolson, “Aerobraking cost and risk decisions,” *Journal of Spacecraft and Rockets*, vol. 44, no. 6, pp. 1285–1293, 2007.
- [7] K. Okamoto, M. Goldshtein, and P. Tsiotras, “Optimal covariance control for stochastic systems under chance constraints,” *IEEE Control Systems Letters*, vol. 2, no. 2,




Figure 8. Evaluation of different cost metrics of the conditional delta-V distribution at 160 seconds after AI.




Figure 9. Comparison of conditional density of delta-V at 160 seconds after AI for different switch times.


pp. 266–271, April 2018.

- [8] J. Ridderhof and P. Tsiotras, “Minimum-fuel powered descent in the presence of random disturbances,” in *2019 AIAA Guidance, Navigation, and Control Conference*, no. AIAA 2019-0646, San Diego, California, 2019.
- [9] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath, “Coherent measures of risk,” *Mathematical Finance*, vol. 9, no. 3, pp. 203–228, 1999.
- [10] Y. Chow, A. Tamar, S. Mannor, and M. Pavone, “Risk-sensitive and robust decision-making: a CVaR optimization approach,” in *Advances in Neural Information Processing Systems*, Montreal, Canada, 2015, pp. 1522–1530.
- [11] E. Tse, “On the optimal control of stochastic linear systems,” *IEEE Transactions on Automatic Control*, vol. 16, no. 6, pp. 776–785, 1971.
- [12] H. A. Simon, “Dynamic programming under uncertainty with a quadratic criterion function,” *Economet-*

rica, Journal of the Econometric Society, vol. 24, no. 1, pp. 74–81, January 1956.

- [13] E. Tse and Y. Bar-Shalom, “Generalized certainty equivalence and dual effect in stochastic control,” *IEEE Transactions on Automatic Control*, vol. 20, no. 6, pp. 817–819, 1975.
- [14] S. Reece and S. Roberts, “An introduction to Gaussian processes for the Kalman filter expert,” in *13th International Conference on Information Fusion*. IEEE, July 2010.
- [15] R. T. Rockafellar, S. Uryasev *et al.*, “Optimization of conditional value-at-risk,” *Journal of Risk*, vol. 2, pp. 21–42, 2000.
- [16] L. Blackmore and M. Ono, “Convex chance constrained predictive control without sampling,” in *AIAA Guidance, Navigation, and Control Conference*, no. AIAA 2009-5876, Chicago, Illinois, August 2009.
- [17] N. Vinh, A. Busemann, and R. Culp, *Hypersonic and Planetary Entry Flight Mechanics*. University of Michigan Press, 1980.
- [18] P. Lu, C. J. Cerimele, M. A. Tigges, and D. A. Matz, “Optimal aerocapture guidance,” *Journal of Guidance, Control, and Dynamics*, vol. 38, no. 4, pp. 553–565, 2015.
- [19] J. M. Laffleur, “The conditional equivalence of ΔV minimization and apoapsis targeting in numerical predictor-corrector aerocapture guidance,” NASA, Technical Report TM-2011-216156, 2011.
- [20] C. Justus, B. James, S. Bougher, A. Bridger, R. Haberle, J. Murphy, and S. Engel, “Mars-GRAM 2000: A Mars atmospheric model for engineering applications,” *Advances in Space Research*, vol. 29, no. 2, pp. 193–202, 2002.
- [21] D. T. Ward and J. Shipley, Buford W., “Control algorithms for aerobraking in the Martian atmosphere,” NASA, Technical Report NASA-CR-185665, 1991.

BIOGRAPHY



Jack Ridderhof is an MS student in mathematics and a PhD student in aerospace engineering at the Georgia Institute of Technology. His research focuses on stochastic control for spacecraft guidance. Jack is a NASA Space Technology Fellow.



Panagiotis Tsiotras is currently the David and Andrew Lewis Endowed Chair Professor at the School of Aerospace Engineering, Georgia Institute of Technology (Georgia Tech), where he is also an Associate Director at the Institute for Robotics and Intelligent Machines. His research interests include optimal control of nonlinear systems and ground, aerial, and space vehicle autonomy. He is a Fellow of AIAA, IEEE, and AAS. He was a recipient of the NSF CAREER Award, the Outstanding Aerospace Engineer Award from Purdue, and the IEEE Award for Technical Excellence in Aerospace Control.